Prepared for submission to JCAP

zBEAMS: A unified solution for supernova cosmology with redshift uncertainties

Ethan Roberts,[1,2] **Michelle Lochner,**[1,4,5] **José Fonseca,**[6] **Bruce A. Bassett,**[1,2,3] **Pierre-Yves Lablanche,**[1] **Shankar Agarwal**[1]

[1]African Institute for Mathematical Sciences,
6 Melrose Road, Muizenberg, 7945, Cape Town, South Africa
[2]Department of Maths and Applied Maths, University of Cape Town,
Cape Town, South Africa
[3]South African Astronomical Observatory,
Observatory, Cape Town, 7925, South Africa
[4]SKA South Africa, 3rd Floor, The Park, Park Road,
Pinelands, 7405, South Africa
[5]Department of Physics and Astronomy, University College London,
Gower Street, London WC1E 6BT, UK
[6]Department of Physics & Astronomy, University of the Western Cape,
Cape Town 7535, South Africa

E-mail: rbreth001@myuct.ac.za

Abstract. Supernova cosmology without spectra will be the bread and butter mode for future surveys such as LSST. This lack of supernova spectra results in uncertainty in the redshifts which, if ignored, leads to significantly biased estimates of cosmological parameters. Here we present a hierarchical Bayesian formalism – zBEAMS – that fully addresses this problem by marginalising over the unknown or contaminated supernova redshifts to produce unbiased cosmological estimates that are competitive with entirely spectroscopic data. zBEAMS provides a unified treatment of both photometric redshifts and host galaxy misidentification (occurring due to chance galaxy alignments or faint hosts), effectively correcting the inevitable contamination in the Hubble diagram. Like its predecessor BEAMS, our formalism also takes care of non-Ia supernova contamination by marginalising over the unknown supernova type. We demonstrate the effectiveness of this technique with simulations of supernovae with photometric redshifts and host galaxy misidentification. A novel feature of the photometric redshift case is the important role played by the redshift distribution of the supernovae.

arXiv:1704.07830v1 [astro-ph.CO] 25 Apr 2017

Contents

1 Introduction **1**

2 Standard Supernova Cosmology Inference **2**

3 Inference in the presence of redshift uncertainties **3**
3.1 Spectroscopic galaxy redshifts but unknown host galaxy identity 5
3.2 Photometric redshifts 6

4 Contamination from non-Ia supernovae **7**

5 The general case: both type and redshift contamination **8**

6 Simulations **10**
6.1 Spectroscopic Case 10
6.2 Photometric Case 11

7 Conclusions and Future Work **14**

1 Introduction

Studies of Type Ia supernovae led to the dark energy breakthrough and modern concordance cosmology, but one can argue that they had been supplanted by Baryon Acoustic Oscillations [1, 2] as the most precise way of constraining cosmology today. To be competitive in the era of LSST [3], EUCLID and the SKA, supernova cosmology faces several big challenges. One is that better control of systematics is required and a number of sophisticated approaches are being developed to improve the control of systematics (e.g. [4–9]). Another critical problem is that next-generation supernova surveys will be severely spectroscopy limited: LSST will deliver over 10^5 Type Ia Supernova (SNIa) candidates with photometric lightcurves only. The lack of spectroscopy introduces a number of challenges. First, the true identity of any candidate without spectroscopic follow-up is ambiguous - photometric colours only provide a probability for an object to be a SNIa, as opposed to a Type Ibc or II supernova or other transient [10]. Secondly, the precise redshifts of the supernovae are unknown. Photometric redshifts are fairly good if the candidates are *known* to be SNIa, yielding RMS errors of $\sigma_z \sim 0.04(1+z)$, depending on exact assumptions [3, 11, 12]. The problem is that we are exactly in the case where we are not sure whether each candidate is a SNIa or not, and the photometric redshift error is much larger if the object is not a SNIa [13], precisely because they are not standard candles.

A promising approach that dates back to the SDSS II supernova survey [14–16], is to obtain spectroscopic redshifts for the host galaxies of the supernova candidates and use this as a proxy for the supernova redshift. This will be particularly attractive in the era of big redshift surveys such as 4MOST, SKA and Euclid, where huge numbers of galaxy redshifts will be known. This has the potential to help remove biases [16] and yield improved constraints [17].

However, even this approach has a serious problem: identifying the host galaxy is also not unambiguous. The supernova can appear to lie in between two or more galaxies or may live in

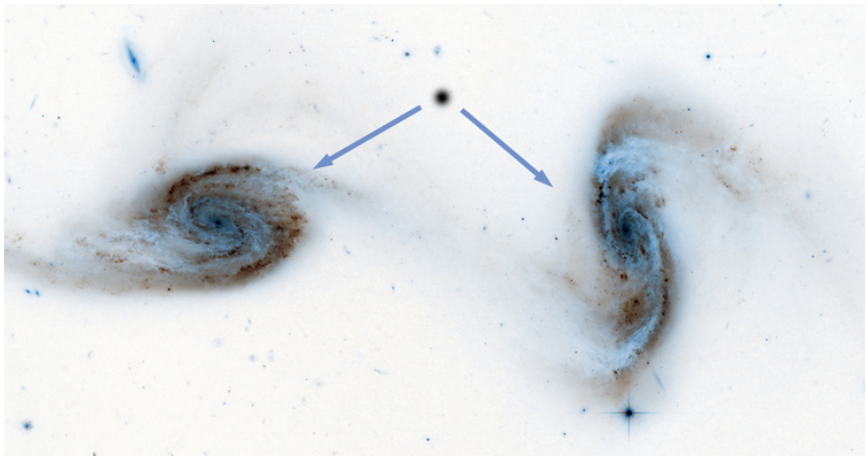

Figure 1. Schematic figure illustrating the source of redshift contamination: the host galaxy of a supernova may be ambiguous even if the redshifts of the galaxies nearby are known. Here we assume that there is an estimate of the probability, denoted $P(\gamma)$, that the supernova belongs to each galaxy (where each galaxy is indexed by γ).

a host that is too faint to be detectable ("hostless") (see figure 1). In general we therefore can assume that instead we have probabilities for the supernova to belong to each of the nearby galaxies on the sky or to be hostless. Current matching algorithms can accurately match the correct host galaxy about 91% of the time when applied to data, potentially increasing to 97% by using machine learning techniques [18]. However, even a 3% contamination may cause significant biases on cosmological parameter inference and must be dealt with.

Our goal with this paper is to develop a single formalism that solves this set of problems, simultaneously handling both the contamination from non-Ia supernovae and contamination from incorrect host identification and redshift assignment. The formalism we desire will be rigorous without resorting to cuts, which are statistically suboptimal.

In laying out the solution we will develop intuition by building systematically to the complete solution. In section 2 we review inference using standard cosmology. In section 3 we consider the case of a single, known Type Ia supernova with a photometric redshift derived from the supernova lightcurve. In section 3.1 we extend the analysis to include host galaxy redshifts. In section 4 we review the BEAMS formalism to handle contamination from non-Ia. In section 5 we present the combined formalism to handle all contamination and finally in section 6 we present an illustrative set of simulations demonstrating our method.

2 Standard Supernova Cosmology Inference

Traditionally, supernova cosmological analysis proceeds with a sample of spectroscopically confirmed type Ia supernovae with well-measured redshifts. The goal is to determine the posterior distribution, $P(\boldsymbol{\theta}|\{D_i\})$, over the cosmological parameters, $\boldsymbol{\theta}$, given the dataset which we denote $\{D_i\}$.[1] Here, the data are the redshift, z_i, and distance modulus, μ_i, for each supernova. The distance modulus for a type Ia supernova is usually estimated from its observed light curve using the SALT2 model [19].

[1] In much of our discussion we will use D to denote a single supernova so for clarity we use $\{D_i\}$ to make explicit the case where we are considering multiple supernovae.

The distance modulus is defined as:

$$\mu(z) = m - M = 5\log_{10}\left(\frac{d_L}{1Mpc}\right) + 25, \tag{2.1}$$

where m is the apparent magnitude of the object, M is the absolute magnitude of the object and d_L is the luminosity distance to the object in Mpc. In a ΛCDM universe, the luminosity distance is related to the cosmological parameters by:

$$d_L(z) = \frac{c(1+z)}{H_0\sqrt{-\Omega_k}}\sin\left(H_0\sqrt{-\Omega_k}\int\frac{dz'}{H(z')}\right), \tag{2.2}$$

where

$$H(z) = H_0\Big(\Omega_m(1+z)^3 + \Omega_\Lambda + \Omega_k(1+z)^2\Big)^{1/2}, \tag{2.3}$$

and H_0 is the Hubble constant, Ω_m is the energy density of matter, Ω_Λ is the energy density of cosmological constant and Ω_k is the curvature parameter. Collectively, we refer to these cosmological parameters as $\boldsymbol{\theta}$.

The posterior probability distribution over the cosmological parameters is then given by Bayes' theorem:

$$P(\boldsymbol{\theta}|\{D_i\}) = \frac{P(\{D_i\}|\boldsymbol{\theta})P(\boldsymbol{\theta})}{P(\{D_i\})}, \tag{2.4}$$

where $P(\{D_i\}|\boldsymbol{\theta})$ is the likelihood, $P(\boldsymbol{\theta})$ is the prior and $P(\{D_i\})$ is called the Bayesian evidence. Since we are not interested in model selection $P(\{D_i\})$ is irrelevant for parameter inference and we drop it in all subsequent analysis.

In the case of uncorrelated Gaussian errors, σ^i_μ, on the estimated distance moduli $\mu^i_{\rm obs}$, the likelihood over the N supernovae is:

$$P(\{D_i\}|\boldsymbol{\theta}) = \prod_i^N \frac{1}{\sqrt{2\pi{\sigma^i_\mu}^2}}\exp\left(-\frac{(\mu^{\rm i}_{\rm obs} - \mu({\rm z^i}, \boldsymbol{\theta}))^2}{2{\sigma^{\rm i}_\mu}^2}\right). \tag{2.5}$$

A standard cosmological analysis would then proceed by applying a numerical sampling method such as MCMC [20, 21] to determine the full posterior of the cosmological parameters. Care needs to be taken around the parameters involved in estimating the distance modulus, such as considered in [22]. While such caveats can be handled easily by introducing new latent parameters, for simplicity we omit these parameters and assume the distance moduli can be measured directly since it is not core to the problem we are addressing.

3 Inference in the presence of redshift uncertainties

To achieve our goal of dealing with the unknown redshifts and contamination we will need two pieces of Bayesian technology in addition to Bayes theorem: marginalisation and the product rule, which we briefly review here. The Product Rule states that, for any sets of parameters $\boldsymbol{\theta}, \boldsymbol{\varphi}$ and $\mathbf{C}$, we have:

$$P(\boldsymbol{\theta}, \boldsymbol{\varphi}|\mathbf{C}) = P(\boldsymbol{\theta}|\boldsymbol{\varphi}, \mathbf{C})P(\boldsymbol{\varphi}|\mathbf{C})\,. \tag{3.1}$$

The second piece of technology is marginalisation over nuisance parameters, $\boldsymbol{\varphi}$:

$$P(\boldsymbol{\theta}|\mathbf{C}) = \int P(\boldsymbol{\theta}, \boldsymbol{\varphi}|\mathbf{C})\; d\boldsymbol{\varphi}\,. \tag{3.2}$$

The combination of these symbol manipulation techniques will be useful to deal with latent parameters; such as the true (unknown) redshift, z, and type τ, of the supernova.

As a step to the fully general case, let us consider a single, known Type Ia supernova with an uncertain redshift, such as that derived from the supernova lightcurve. Uncertainties in the redshift, δz, have usually been dealt with by converting them into an additional error in the distance modulus, $\delta\mu$, by assuming a model $\mu(z)$ for the conversion, and added the result in quadrature with the usual $\mu_{\rm obs}$ error, but this is not statistically correct (for one since we don't know the correct distance modulus describing our universe) as we will discuss in detail later.

To proceed rigorously we instead start by using Bayes theorem for the posterior for $\boldsymbol{\theta}$ given some data, D, from a single supernova. We then expand the arguments of the posterior to include the true redshift (z) and distance modulus (μ) of the supernova as latent (i.e. nuisance) parameters that we marginalise over, since we don't know their true values:

$$P(\boldsymbol{\theta}|D) \propto P(D|\boldsymbol{\theta})P(\boldsymbol{\theta})\,, \tag{3.3}$$

$$\propto \int P(D,z,\mu|\boldsymbol{\theta})P(\boldsymbol{\theta})\,dz\,d\mu\,. \tag{3.4}$$

We repeatedly apply the product rule to rewrite this multi-dimensional integral as:

$$P(\boldsymbol{\theta}|D) \propto \int P(D|z,\mu,\boldsymbol{\theta})P(z,\mu|\boldsymbol{\theta})P(\boldsymbol{\theta})\,dz\,d\mu\,, \tag{3.5}$$

$$\propto \int P(D|z,\mu)P(\mu|z,\boldsymbol{\theta})P(z|\boldsymbol{\theta})P(\boldsymbol{\theta})\,dz\,d\mu\,. \tag{3.6}$$

We now make the simplifying assumption that the distribution of the true redshift is independent of the cosmological parameters[2] and note that since μ is assumed to be a deterministic function of z and $\boldsymbol{\theta}$ the distribution $P(\mu|z,\boldsymbol{\theta})$ is a delta function[3], which allows us to eliminate the μ-integral[4]:

$$P(\boldsymbol{\theta}|D) \propto \int P(D|z,\mu)\,\delta(\mu-\mu(z,\boldsymbol{\theta}))P(z)P(\boldsymbol{\theta})\,dz\,d\mu\,, \tag{3.7}$$

$$\propto \int P\left(D|z,\mu(z,\boldsymbol{\theta})\right)\,P(z)\,P(\boldsymbol{\theta})\,dz\,. \tag{3.8}$$

This is the expression we were seeking. It expresses the posterior as a marginalisation over the unknown supernova redshift. What is the data D? In this case let us assume that we have extracted both an estimate of the redshift and distance modulus, $(z_{\rm obs},\mu_{\rm obs})$, for the supernova from its lightcurve. We now have:

$$P(\boldsymbol{\theta}|D) \propto P(\boldsymbol{\theta})\int_0^\infty P\left(z_{\rm obs},\mu_{\rm obs}|z,\mu(z,\boldsymbol{\theta})\right)\,P(z)\,dz\,. \tag{3.9}$$

Considering that the estimates $z_{\rm obs}$ and $\mu_{\rm obs}$ are independent, we can simplify this to:

$$P(\boldsymbol{\theta}|D) \propto P(\boldsymbol{\theta})\int_0^\infty P\left(z_{\rm obs}|z\right)P\left(\mu_{\rm obs}|\,\mu(z,\boldsymbol{\theta})\right)\,P(z)\,dz\,. \tag{3.10}$$

[2] A more accurate assumption would be to posit that it depends on $\boldsymbol{\theta}$ only via the volume of spacetime. This is expected to be a weak dependence for the currently allowed range for $\boldsymbol{\theta}$.

[3] This is only exactly true in the background FLRW model and is not true if one allows for effects such as gravitational lensing, but we will ignore such perturbative effects here as is typical in supernova studies.

[4] Using the standard identity $\int f(x)\delta(x-x_0)dx = f(x_0)$

An important novel role in this analysis is played by the redshift prior $P(z)$. This has no parallel in the usual supernova analysis where the redshift is known spectroscopically. In the case of the uncertain supernova redshift we must consider Eddington bias: a supernova discovered with a 4m telescope with a redshift estimate of $z_{\rm obs} = 0.75$, is much more likely to lie at a *true* redshift of $z = 0.6$ than at $z = 0.9$.

This insight, encoded in the prior $P(z)$, must be included in our inference analysis. Gull [23] was perhaps the first to highlight the importance of the data prior, showing that the use of the wrong prior (i.e. not the true distribution from which the data was drawn) leads to a bias.

Including this is crucial in any practical application since we are not sure about the true redshift prior, $P(z)$, which will depend both on the cosmological volume (which depends on the cosmological parameters) and the rates for each type of supernova. We can handle this uncertainty by introducing hyperparameters, φ, into the prior, $P(z, \varphi)$ which can then be fit for or marginalised over. Indeed, one can trivially turn this around and instead marginalise over the cosmological parameters and supernova redshifts to produce a posterior for the rates of the different supernova types as functions of redshift. Hence this "annoyance" which cannot be handled by the standard formalism becomes a feature in our analysis.

While this will be important for realistic simulations and analysis of real data it is a conceptually straightforward Bayesian extension and we do not consider it further here.

We now consider two special limiting cases to gain some intuition: (i) Spectroscopic galaxy redshifts but unknown host identity, and (ii) photometric supernova redshifts alone.

3.1 Spectroscopic galaxy redshifts but unknown host galaxy identity

In this subsection we answer the question: what happens if we have spectroscopic redshifts for all galaxies but are not sure which galaxy is the true host?

Since we take our data to be coming from the supernova lightcurve only, we can treat spectroscopic galaxy redshifts as a prior $P(z)$ on the SN redshift. Since we don't know the true host in general we write this prior as a marginalisation over the potential host galaxies, which we index by γ:

$$P(z) = \sum_\gamma P(z|\gamma)P(\gamma)\,. \tag{3.11}$$

In the case of figure 1 this sum would likely just consist of just three terms: the two nearby galaxies and a third invisible galaxy which might be too faint to be detected in the image.

In eq. (3.11) $P(\gamma)$ is the probability that the supernova is hosted by the $\gamma^{\rm th}$ galaxy[5], and $P(z|\gamma)$ is the corresponding redshift distribution for the $\gamma^{\rm th}$ galaxy. Then we have:

$$P(\boldsymbol{\theta}|D) \propto P(\boldsymbol{\theta}) \int dz \sum_\gamma P(z|\gamma)P(\gamma)P\left(z_{\rm obs}, \mu_{\rm obs}|z, \mu(z, \boldsymbol{\theta})\right) \tag{3.12}$$

In the limit where the galaxy redshifts are known spectroscopically to high precision we can approximate the host redshift distribution as a delta function: $P(z|\gamma) = \delta(z - z_\gamma)$ where z_γ is the spectroscopically determined redshift of the $\gamma^{\rm th}$ galaxy. Then we can perform the marginalisation analytically, yielding:

$$P(\boldsymbol{\theta}|D) \propto P(\boldsymbol{\theta}) \sum_\gamma P(\gamma)P\left(z_{\rm obs}, \mu_{\rm obs}|z_\gamma, \mu(z_\gamma, \boldsymbol{\theta})\right)\,. \tag{3.13}$$

[5]Estimated for example using the projected distance from the centre of each galaxy of some other measure.

In other words, the final posterior is a mixture model that simply sums all the posteriors arising from assuming the supernova belongs to each of the potential host galaxies, weighted by the probability $P(\gamma)$ that the supernova belongs to each host. This is an intuitively pleasing and simple result.

3.2 Photometric redshifts

A second limiting subcase that will give some useful intuition is the case where we have no spectroscopic host information but instead only have a photometric redshift estimate from the supernova itself or from the host galaxy in the case where the host is unambiguous, i.e. we have an estimate for $P\left(z_{\rm obs}|z\right)$.

For simplicity assume that the resulting photometric redshift distribution is Gaussian[6] $P\left(z_{\rm obs}|z\right) = \exp(-(z-z_{\rm obs})^2/2\sigma_z^2)/\sqrt{2\pi\sigma_z^2}$. Assuming independent $z_{\rm obs}$ and $\mu_{\rm obs}$ estimates for simplicity, eq. (3.9), reduces to:

$$P(\boldsymbol{\theta}|D) \propto P(\boldsymbol{\theta}) \int \frac{1}{2\pi\sigma_z\sigma_\mu} \exp\left(-\frac{(z_{\rm obs}-z)^2}{2\sigma_z^2}\right) \exp\left(-\frac{(\mu_{\rm obs}-\mu(z,\boldsymbol{\theta}))^2}{2\sigma_\mu^2}\right) P(z)dz \quad (3.14)$$

In our simulations we numerically perform the marginalisation in eq. (3.14) without considering correlations between $z_{\rm obs}$ and $\mu_{\rm obs}$. Such correlations can be included in a straight forward way via modeling the covariance function, potentially with hyperparameters which can also be marginalised over.

Examining eq. (3.14) one might be tempted to Taylor expand $\mu(z,\boldsymbol{\theta})$ around $z_{\rm obs}$. Taking only the linear term in $\Delta z \equiv (z - z_{\rm obs})$ one can, assuming a Gaussian prior $P(z)$ with mean $\bar{z}$ and standard deviation σ_p, perform the redshift marginalisation analytically [7], giving:

$$P(\boldsymbol{\theta}|D) \propto P(\boldsymbol{\theta}) \int \exp\left(-\frac{(z_{\rm obs}-z)^2}{2\sigma_z^2} - \frac{(\tilde{\mu}_{\rm obs}-\mu' z)^2}{2\sigma_\mu^2}\right) \exp\left(-\frac{(z-\bar{z})^2}{2\sigma_p^2}\right) dz\,, \quad (3.15)$$

where $\tilde{\mu}_{\rm obs} = \mu_{\rm obs} - \mu(z_{\rm obs}) + \mu' z_{\rm obs}$. Completing the square and performing the integral gives:

$$P(\boldsymbol{\theta}|D) \propto \frac{P(\boldsymbol{\theta})}{C^{1/2}} \exp\left(-\frac{1}{2C}\left[\sigma_\mu^2(z_{\rm obs}-\bar{z})^2 + \sigma_p^2(\mu_{\rm obs}-\mu(z_{\rm obs}))^2 + \sigma_z^2(\tilde{\mu}_{\rm obs}-\mu'\bar{z})^2\right]\right)\,, \quad (3.16)$$

where C is equal to :

$$C = \sigma_\mu^2(\sigma_p^2+\sigma_z^2) + \mu'^2\sigma_z^2\sigma_p^2\,. \quad (3.17)$$

This has a simple interpretation: the z-error is converted into a μ error by using μ', the instantaneous slope of the $\mu(z,\boldsymbol{\theta})$ curve, which is then added in quadrature with the existing μ error, σ_μ^2; all modulated by the redshift prior width σ_p. Since we are assuming a Gaussian likelihood this is the same result as for the corresponding Fisher matrix analysis, see [24]. Note that to do this self-consistently one must change the μ error for each supernova at every point in the MCMC chain (since changing $\boldsymbol{\theta}$ changes μ').

On the surface this offers a big computational simplification. For a large dataset of N supernovae and m cosmological parameters, eq. (3.16) only requires one to fit for the m parameters rather than all $N+m$ parameters. Unfortunately in most cases eq. (3.16) is not sufficient for several reasons:

[6] The generalisation to an arbitrary distribution is trivial since we perform the marginalisation numerically.

[7] Note that one cannot simply assume an improper uniform prior since this leads to biases with real data while assuming a proper top-hat prior is no longer analytically integrable.

- It is only applicable when $P(z_{\rm obs}|z)$ is sufficiently narrow, i.e. $\sigma_\mu^2 \gg \sigma_z^2 \mu'^2$. But what does this mean in practise?
- As alluded to previously we can only perform the integral in eq. (3.14) analytically for special priors such as an improper uniform ($-\infty < z < \infty$) or Gaussian prior.[8] These are not appropriate for astronomical surveys and hence assuming one of them will give biased results in general [23].
- In the general case we do not know the type of the supernova for sure. This implies that the photometric redshift for the supernova may be very non-Gaussian in general. This in turn implies that we cannot analytically integrate eq. (3.14) even if we assume a Gaussian redshift prior.

In the simulations below we find that this approach is biased. For all these reasons we strongly recommend numerically marginalising over all redshifts via MCMC or suitable sampling technique as we present in subsection 3.2.

4 Contamination from non-Ia supernovae

In the previous section we derived the posterior in the presence of redshift contamination but assumed that we knew the object was a SNIa. Unfortunately contamination of supernova types is also inevitable for photometric surveys. This has been addressed by the BEAMS (Bayesian Estimation Applied to Multiple Species) formalism for the case of spectroscopic redshifts in various papers [5, 14, 25, 26]. Here we highlight the key results using the hierarchical Bayesian approach we have used so far in this paper for redshift contamination as a warm-up to the case with both types of contamination.

In this case we can assume we know the true redshift of the supernova but are not sure of its type, labeled by a discrete variable τ, which we here allow to take two values: $\tau = \text{Ia}$ and $\tau = \not{\text{Ia}}$, the latter corresponding to non-Ia objects.[9]

Previously we marginalised over the unknown latent variables z and μ. Now instead we marginalise over τ, μ and z. Following the same approach taken in eq. (3.6) and thereafter, i.e., using Bayes theorem and the Product Rule to marginalise over the latent variables τ, μ and z, one has

$$P(\boldsymbol{\theta}|D) \propto P(D|\boldsymbol{\theta})P(\boldsymbol{\theta})\,, \tag{4.1}$$

$$\propto \int P(D|z,\mu,\tau,\boldsymbol{\theta})P(z,\mu|\tau,\boldsymbol{\theta})P(\tau)P(\boldsymbol{\theta})\,d\tau\,dz\,d\mu\,, \tag{4.2}$$

$$\propto \int P(D|z,\mu,\tau,\boldsymbol{\theta})P(\mu|z,\tau,\boldsymbol{\theta})P(z|\tau,\boldsymbol{\theta})P(\tau)P(\boldsymbol{\theta})\,d\tau\,dz\,d\mu\,. \tag{4.3}$$

Since we assume we know the redshift of the supernova, z_*, perfectly, we can write $P(z|\tau,\boldsymbol{\theta})$ as a delta function $\delta(z - z_*)$ which allow us to do the the z-integral, yielding:

$$P(\boldsymbol{\theta}|D) \propto \int P(D|z_*,\mu,\tau,\boldsymbol{\theta})P(\mu|z_*,\tau,\boldsymbol{\theta})P(\tau)P(\boldsymbol{\theta})\,d\tau\,d\mu\,. \tag{4.4}$$

[8] This is integrable for a prior $P(z) \propto ze^{-\beta z}$ but this doesn't have enough freedom in general for our needs.

[9] If there were more than one class of object that gave useful information about $\boldsymbol{\theta}$ we would need to subdivide the non-Ias into more classes; e.g. if the $\boldsymbol{\theta}$ represented information about star formation rates rather than cosmology.

As before, μ is a deterministic function of z, τ and $\boldsymbol{\theta}$ and hence $P(\mu|z_*, \tau, \boldsymbol{\theta})$ is also a delta-function, which collapses the μ-integral. We write $\mu_{\rm Ia}$ for the expression $\mu(z, \boldsymbol{\theta}, \tau = {\rm Ia})$ and $\mu_{\not{\rm Ia}}$ for $\mu(z, \boldsymbol{\theta}, \tau = \not{\rm Ia})$. The τ 'integral' is actually just a sum over the two supernova classes. Here, $P(\tau)$ is the supernova type probability.[10] For notational simplicity we write the probability of being a type Ia supernova as $P_{\rm Ia}$ (as for $P(\tau = {\rm Ia})$) and the probability of not being a type Ia supernova as $P_{\not{\rm Ia}} = 1 - P_{\rm Ia}$.

Then eq. (4.4) reduces to the usual BEAMS result:

$$P(\boldsymbol{\theta}|D) \propto P(\boldsymbol{\theta}) \left[P_{\rm Ia}\, P(D|z_*, \mu_{\rm Ia}(z), \tau = {\rm Ia}, \boldsymbol{\theta}) + (1 - {\rm P_{Ia}})\, {\rm P(D|z_*, \mu_{\not Ia}(z), \tau = \not Ia, \boldsymbol{\theta})}\right] . \quad (4.5)$$

For a Gaussian likelihood with observed distance modulus $\mu_{\rm obs}$, we can write:

$$P(D|z_*, \mu_{\rm Ia}(z), \tau = {\rm Ia}, \boldsymbol{\theta}) = \frac{1}{\sqrt{2\pi}\sigma_{\rm Ia}} \exp\left(\frac{-(\mu_{\rm obs} - \mu_{\rm Ia}({\rm z_*}, \boldsymbol{\theta}))^2}{2\sigma_{\rm Ia}^2}\right) . \quad (4.6)$$

What about the non-Ia term $\mu_{\not{\rm Ia}}$? Following [14] and [25] one can take $\mu_{\not{\rm Ia}} = \mu_{\rm Ia} + b(z, \varphi)$ where $b(z, \varphi)$ allows both for the redshift evolution of the non-Ia population and the difference in mean intrinsic luminosity compared with SNIa, parameterised by some unknown hyper-parameters φ which are fit simultaneously with all the other parameters. This allows one to learn about the non-Ia populations. For cosmology, the non-Ia population is essentially uninformative and hence one can simply take the corresponding $\sigma_{\not{\rm Ia}}$ to be large in actual analysis.

5 The general case: both type and redshift contamination

We are now ready to consider the general case with both type and redshift contamination. Hence we assume that the type τ of the supernova is unknown, and that we either have photometric redshift estimate for the supernova or redshift information (either photometric or spectroscopic) of potential hosts galaxies.

As before we will treat the host galaxy information as a prior, $P(z)$, on the SN redshift (see eq. (3.11))[11]:

$$P(z) = \sum_{\gamma} P(z|\gamma) P(\gamma) . \quad (5.1)$$

These need not be spectroscopic redshifts and $P(z|\gamma)$ may have significant spread.

As before we begin from Bayes rule, introduce and marginalise over latent variables for

[10] This probability is estimated directly from the lightcurve using templates, inference or machine learning [10, 27].

[11] The supernova rates are now a much less important factor in the prior since the redshift priors arising from the galaxies should be very peaked in comparison.

the unknown true variables, z, μ, τ, and then repeatedly apply the product rule:

$$P(\boldsymbol{\theta}|D) \propto \int P(D, \tau, z, \mu|\boldsymbol{\theta}) P(\boldsymbol{\theta})\, d\tau\, dz\, d\mu\,, \tag{5.2}$$

$$\propto \int P(D|\tau, z, \mu, \boldsymbol{\theta}) P(\tau, z, \mu|\boldsymbol{\theta}) P(\boldsymbol{\theta})\, d\tau\, dz\, d\mu\,, \tag{5.3}$$

$$\propto \int P(D|\tau, z, \mu) P(\mu|\tau, z, \boldsymbol{\theta}) P(\tau, z|\boldsymbol{\theta}) P(\boldsymbol{\theta})\, d\tau\, dz\, d\mu\,, \tag{5.4}$$

$$\propto \int P(D|\tau, z, \mu) P(\mu|\tau, z, \boldsymbol{\theta}) P(\tau|z, \boldsymbol{\theta}) P(z|\boldsymbol{\theta}) P(\boldsymbol{\theta})\, d\tau\, dz\, d\mu\,, \tag{5.5}$$

$$\propto \int P(D|\tau, z, \mu) \delta(\mu - \mu(\tau, z, \boldsymbol{\theta})) P(\tau|z) P(z) P(\boldsymbol{\theta})\, d\tau\, dz\, d\mu\,, \tag{5.6}$$

$$\propto \sum_{\tau} \int P\left(D|\tau, z, \mu(\tau, z, \boldsymbol{\theta})\right) P(\tau|z) P(z) P(\boldsymbol{\theta})\, dz\,, \tag{5.7}$$

where the last step accounts for the fact that τ is a discrete (categorical) variable. We can now substitute for the prior $P(z)$ expressed as the sum over the potential host galaxies, eq. (3.11), which gives:

$$P(\boldsymbol{\theta}|D) \propto P(\boldsymbol{\theta}) \int dz \sum_{\tau} P(\tau|z) P\left(D|\tau, z, \mu(\tau, z, \boldsymbol{\theta})\right) \sum_{\gamma} P(z|\gamma) P(\gamma)\,. \tag{5.8}$$

Eq. (5.8) is our main result for the posterior arising from a single supernova. It handles both Type and redshift-host contamination. Note that because the true redshift of the supernova is unknown we can allow for redshift-dependence of the type probabilities, $P(\tau|z)$, whereas in previous BEAMS analysis this was taken as a constant.

To compute the full posterior over N supernovae with collective data $\{D_i\}$, we now have $2N$ nuisance parameters corresponding to the type and redshift, (τ, z) for each supernova, in addition to the cosmological parameters $\boldsymbol{\theta}$. We can then either assume the supernovae are independent, and simply multiply the posteriors to get:

$$P(\boldsymbol{\theta}|\{D_i\}) \propto P(\boldsymbol{\theta}) \prod_{i}^{N} \left[\int dz \sum_{\tau_i} P(\tau_i|z_i) P\left(D_i|\tau_i, z_i, \mu(\tau_i, z_i, \boldsymbol{\theta})\right) \sum_{\gamma_i} P(z_i|\gamma_i) P(\gamma_i) \right], \tag{5.9}$$

or allow for correlations between the supernovae as in [26].

In practice the marginalisation over redshift required by eq. (5.9) can be achieved efficiently through MCMC by allowing the redshift of each supernova to be a free nuisance parameter that is varied along with the cosmological parameters $\boldsymbol{\theta}$. If the supernovae are correlated then we must also introduce the type of each supernova to include as a parameter. This does not significantly alter the MCMC analysis [26].

At this point one can again ask what the data, D_i, is for each supernova? At the most basic – and correct – level this would be the lightcurve measurements in various colour bands as a function of time. However, a convenient simplification is to consider $D_i = (z_{\rm obs}, \mu_{\rm obs})$, *assuming* that the object is a SNIa. In general this would be inappropriate if more than one type of supernova contained useful information about the cosmological parameters $\boldsymbol{\theta}$. However, since in the non-Ia case the derived redshift and distance modulus give almost no useful cosmological information, one can simply take $P(z_{\rm obs}|z, \tau = \cancel{\rm Ia})$ to be a wide uniform distribution or very wide Gaussian and use $z_{\rm obs}, \mu_{\rm obs}$ derived assuming the object is a SNIa

[5, 25]. Then the fact that one is using the "wrong" values for $z_{\rm obs}, \mu_{\rm obs}$ has no impact. We note that obviously this is only an issue when the type of the supernova is unknown.

We now move to test this formalism in different cases.

6 Simulations

Here we perform an illustrative set of simulations to show how zBEAMS, as described in sections 3 to 5, recovers the correct cosmological parameters by marginalising over unknown supernova types and redshifts. We consider the two cases corresponding to sections 5 and 3.2: one where we have spectroscopic redshift estimates from potential host galaxies, hereafter referred to as the *spectroscopic case*, and one where we have photometric measurements only, hereafter referred to as the *photometric case*. Conceptually, it is trivial to combine these cases when dealing with a dataset with mixed spectroscopic and photometric measurements, however we keep the two cases distinct for simplicity. We demonstrate that while the standard likelihood, ignoring redshift uncertainty, results in biased cosmological parameter estimates, zBEAMS correctly marginalises over type and redshift uncertainties, recovering the fiducial cosmology.

For all simulations, we assume a flat ΛCDM universe with a fiducial cosmology given by the latest results from the Planck collaboration [28], i.e., $H_0 = 67.74$ km/s/Mpc, $\Omega_m = 0.31$ and $w = -1$. We perform inference over the parameters using Markov Chain Monte Carlo (MCMC) methods, specifically we use the Metropolis-Hastings [20, 21] algorithm for low-dimensional sampling and block Metropolis-Hastings when numerically marginalising over redshift in the photometric case. Detailed descriptions of both cases follow below.

6.1 Spectroscopic Case

For the spectroscopic case, we simulate 1000 SNe from a uniform distribution across a redshift range of $z = 0.015 - 1$. We create two datasets, one pure Ia dataset without any redshift errors for reference and one with both host galaxy redshift and non-Ia contamination.

The unbiased SNe data is generated with a dispersion of 0.2 mag. For the biased dataset, we assume the host galaxies have spectroscopically confirmed redshifts, but the supernova is observed using photometry, hence the supernova type is not known and it is not always clear which galaxy the supernova belongs to if multiple galaxies lie within a small angular distance of one another. We also considered that supernovas with $z < 0.1$ will be identified spectroscopically. We assume 5% type misidentification where the non-Ia population is offset from the Ia population by 2 mag, and has a Gaussian dispersion of 1.5 mag (similar to [14]). A more realistic distribution can be used with the same method, provided the form is known. We assume a 9% host misidentification [18], where the misidentified host redshift is drawn from a normal distribution $z \sim \mathcal{N}(z_{\rm true}, 0.1^2)$.

In figure 2 we plot the distance modulus of a contaminated SNe Ia dataset. The second dataset we analyse using both the standard likelihood (which does not take redshift error into account) and with the zBEAMS likelihood. We use the zBEAMS posterior in eq. (5.9) to fully marginalise over both type and redshift uncertainties and thus produce unbiased cosmological estimates. In this analysis, we solve for Ω_m, H_0 and w while we assume the parameters of the populations (such as the magnitude offset and standard deviation of the non-Ia population) are known. However, it would be trivial to solve for these simultaneously. We infer the marginalised posterior distribution for w and Ω_m for each of these three instances, and their respective contours are shown in figure 3.

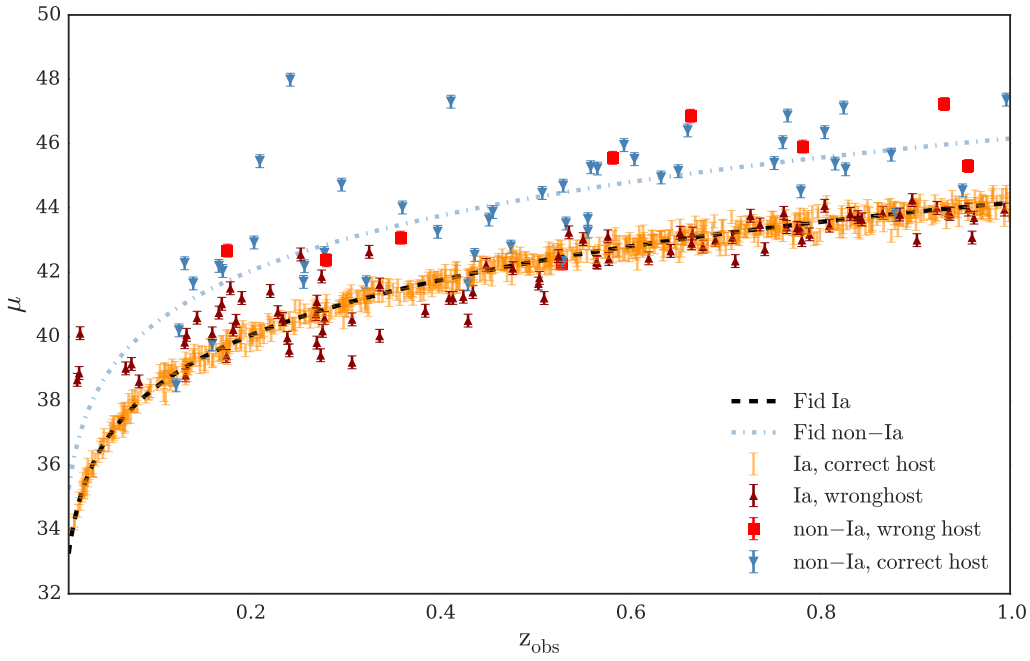

Figure 2. Distance moduli for 1000 SNe with two types of contamination: $\sim 9\%$ wrong host-galaxy identification (and hence incorrect redshift, shown as red and maroon points) and $\sim$ 5% non-Ia contaminants (shown as blue and red data points). The red data points represent SNe that have both the wrong redshift and are non-Ia. For cases with the wrong host galaxy redshift we assign a Gaussian error with a conservative standard deviation of 0.1. The fiducial Ia and non-Ia distance modulus are shown as the thick and thin dashed lines respectively. Figure 3 shows how zBEAMS is able to untangle both forms of contamination with little increase in error contour size, while applying the standard MCMC approach ignoring the contamination leads to very significant biases.

6.2 Photometric Case

Here we consider the case where the redshift of the host galaxy is obtained photometrically. Now to use zBEAMS, the marginalisation over redshift must be performed numerically, as the integral in eq. (3.14) has no analytic solution. We assume for this work that the redshift uncertainties are Gaussian distributed with a standard deviation of $0.04(1+z)$. We simulate 998 SNe from a redshift distribution given by $P(z) \sim ze^{-\beta z}$ spanning the redshift range $z = 0.015 - 1.4$, with $\beta = 3$. Note that this distribution is also the prior on the fitted redshifts and in reality one would need to fit for β simultaneously with all other parameters. While we do not address this here, this distribution could be extended to include modeling of instrumental selection effects, in addition to intrinsic supernova rate information. As before, we assumed Gaussian errors with dispersion 0.2 mag in the distance modulus. In figure 4 one can see the magnitude residuals for the observational redshift (main figure) and for the redshifts recovered by the zBEAMS analysis (inset figure).

We then do a MCMC using the standard likelihood (where we only solve for 3 cosmological parameters assuming incorrectly that all the redshifts are correct). The result is plotted in figure 6 in shades of red. Note that the result is clearly biased with respect to the solid black contours which are obtained using the true SNe redshifts. The blue contours in figure 6 show the result when applying zBEAMS to the biased dataset. We used a block

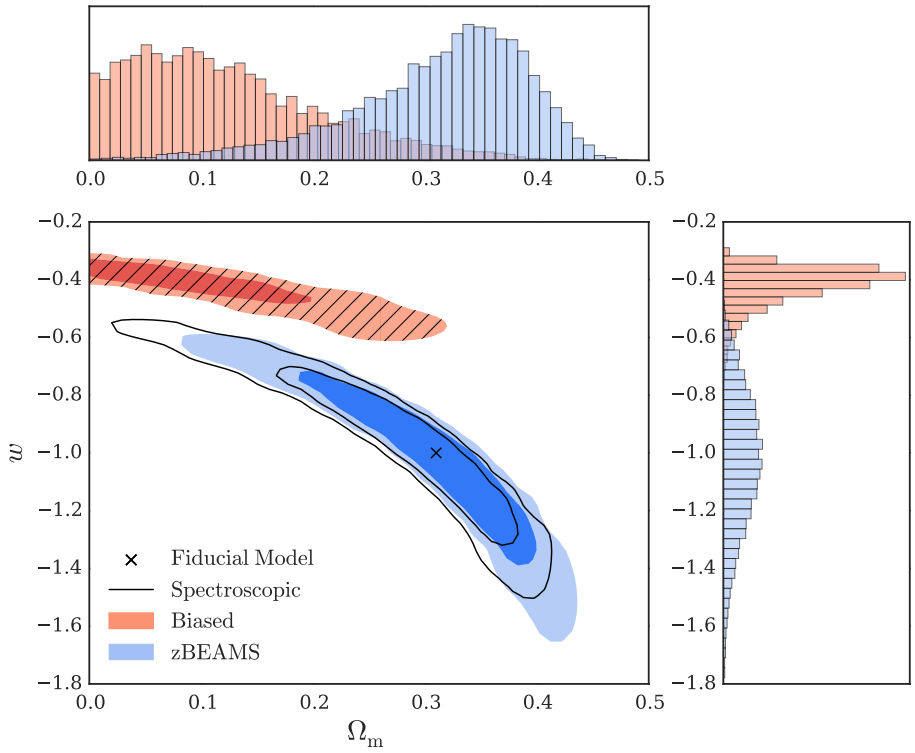

Figure 3. Contour plots for w and Ω_m showing the 68% and 95% credible intervals for the three instances we consider in the spectroscopic case. The black cross shows the fiducial model from which the data were generated. The black lined contours show the ideal "model answer" posterior where we use the standard likelihood on the spectroscopic dataset with no redshift or type contamination. The red solid contours with dashed lines illustrate the posterior when applying standard MCMC without accounting for the 9% redshift contamination with Gaussian redshift confusion ($\sigma_z \sim 0.1$) and 5% non-Ia type confusion. Finally the blue solid contours show the zBEAMS posterior on the same doubly-contaminated dataset. zBEAMS handles both the redshift and type contamination with little increase in computational complexity or error ellipse area. Top and right panels show the 1D marginalised histograms for Ω_m and w respectively for the standard likelihood (red) and the zBEAMS likelihood (blue).

Metropolis-Hastings sampling method – affectionately dubbed "Arabian nights" – to fit for 1001 parameters simultaneously (3 cosmological parameters and 998 redshifts), i.e., numerically computing the posterior given by eq. (3.14). The block Metropolis-Hastings proceeds identically to the usual Metropolis-Hastings sampling algorithm, except that parameters are updated in blocks instead of updating all parameters every step. We took the 3 cosmological parameters and each redshift as a block, but experimented with block sizes of 1-10 redshifts.

The block size has little impact on accuracy, as long as the blocks are small enough not to reduce the acceptance ratio, they do however impact on algorithm speed. We assume that each supernova redshift has a prior coming from the host galaxy (or from the supernova lightcurve itself) which we take to be Gaussian centred on the observed redshift with standard deviation of $0.04(1 + z_{\rm obs})$. The prior on the overall SNIa redshift distribution was taken to

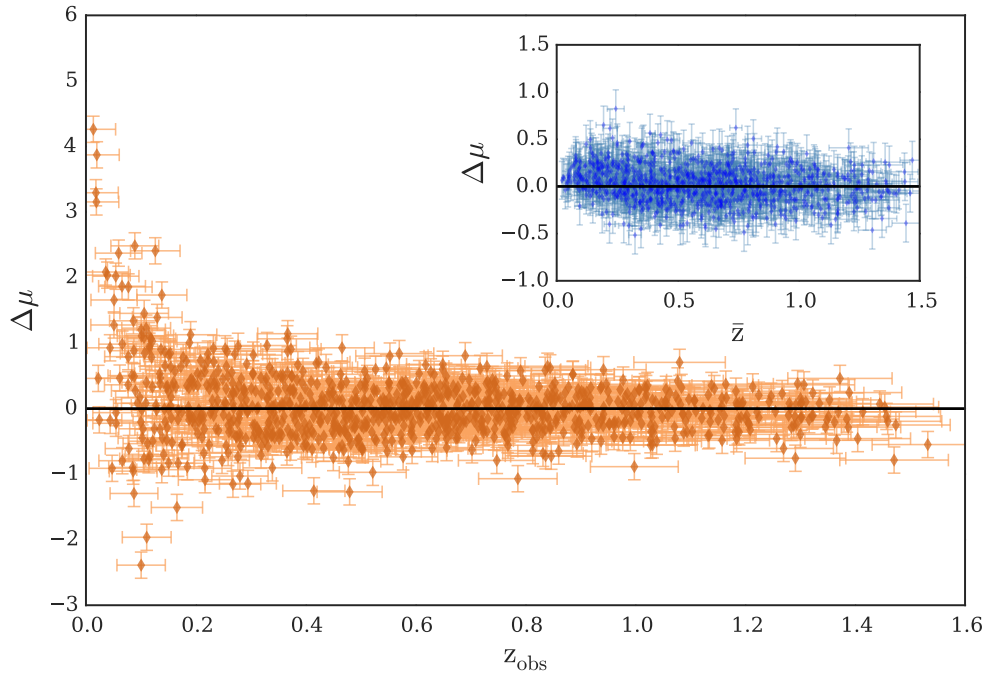

Figure 4. Photometric Hubble residuals for 998 SNIa drawn from the redshift prior distribution $P(z) \sim ze^{-3z}$ with photometric redshift errors drawn from a Gaussian with $\sigma_z = 0.04(1+z)$. Main figure: residuals plotted against $z_{\rm obs}$ (gold). Inset: residuals plotted against the redshifts recovered from the MCMC chains, $\bar{z}$, (blue). The redshift uncertainties cause a large fraction of the data to appear more than 3σ away from the fiducial model, with some points more than 20σ away (the error on μ is 0.2 mag in all cases). Instead zBEAMS handles these large excursions by effectively putting the supernovae at the correct redshift (as shown in the inset where the majority of points are less than 2σ from the fiducial model) and recovers the unbiased cosmology contours shown in figure 6 without significant decrease in precision.

be $P(z) = ze^{-\beta z}$, where we fixed the value of β to 3. In a case with real data, one would need to fit for these hyperparameters as well.

We found that the block Metropolis-Hastings recovers the true redshifts of the low-z supernovae ($\sigma_z = 0.02$ for $z < 0.25$) well, with worsening performance as the redshift increases. This is due to two effects: first we assume the photometric redshift error scales with $(1+z)$ and secondly the Hubble diagram progressively flattens out at $z > 0.25$ removing the signal that allows MCMC to constrain the redshift. This can be clearly seen in figure 5 where we show the stacked 1d histograms $z_{i,\rm chain} - z_{i,\rm true}$ for all 998 SNIa. It can be seen that while the error increases with redshift, the redshift estimates show no systematic bias. The marginalised posterior distributions for w and Ω_m for each of these instances are represented in the contour plots shown in figure 6. We can see that zBEAMS recovers the correct cosmology, and contours, as desired.

Examining figure 4 we can see the origin of the bias of the standard likelihood. A large number of the data points are more than 3σ away from the model and some are over 20σ away. This is an artefact of using the wrong redshifts. The inset shows the same residual Hubble diagram when the data is instead plotted using the means redshifts recovered as the

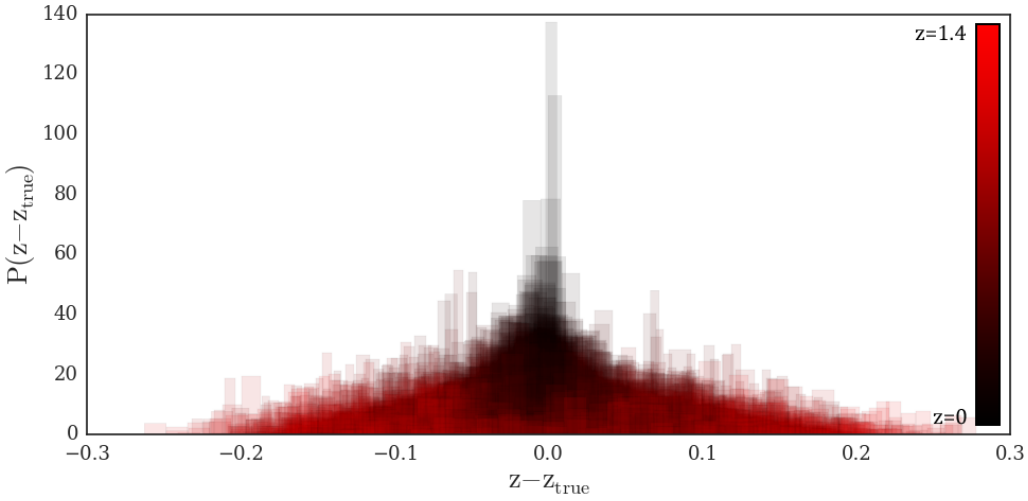

Figure 5. Stacked one-dimensional histograms for all 998 redshifts from the zBEAMS analysis of the data in figure 4. For each supernova we plot the histogram relative to the true redshift, demonstrating that zBEAMS recovers, on average, the true redshift for each supernova. Each histogram is coloured by its redshift: black corresponding to low redshifts and red corresponding to high redshifts, showing that the recovered redshifts are less precise for increasing redshift, as expected due to the $(1 + z)$ scaling of the photometric redshift error and the flattening of the Hubble diagram.

mean of MCMC chain for each redshift. Very few datapoints are now more than 2σ from the fiducial model, especially at low redshifts where the excusions were the strongest, showing how effectively zBEAMS uses the gradient in the distance modulus to recover the true redshifts.

This allows zBEAMS to produce unbiased cosmology contours that almost match the contour sizes of the perfect, spectroscopic case. A more standard approach to this same problem might be to significantly increase all the μ-error bars of the points to account for the redshift uncertainties. We find that doing so yields biased results for reasons discussed in section 3.2. Increasing the error bars further might unbias the contours but only at the expense of significantly inflating the associated contours.

It should be noted that accurate sampling in realistic scenarios will not be trivial since we are fitting for a posterior that typically has more unknown parameters than data points. We have explored both block Metropolis-Hastings and other algorithms such as Diffusive Nested Sampling [29] as viable solutions. Hamiltonian Monte Carlo [30] may also be well-suited to the high-dimensionality of this problem.

The code used in this section is available on Github at:
https://github.com/MichelleLochner/zBEAMS.

7 Conclusions and Future Work

Future large surveys such as LSST will likely deliver large numbers ($> 10^5$) of good supernova candidates without the spectroscopic confirmation that has historically been required to use them for cosmology. This means we will have to do cosmology with a sample for which both the true type and redshift of the supernovae are unknown. Instead, only the probability distributions of both supernova type and redshift will be available. In particular, the redshifts will be uncertain either because only a redshift is only known photometrically or because the

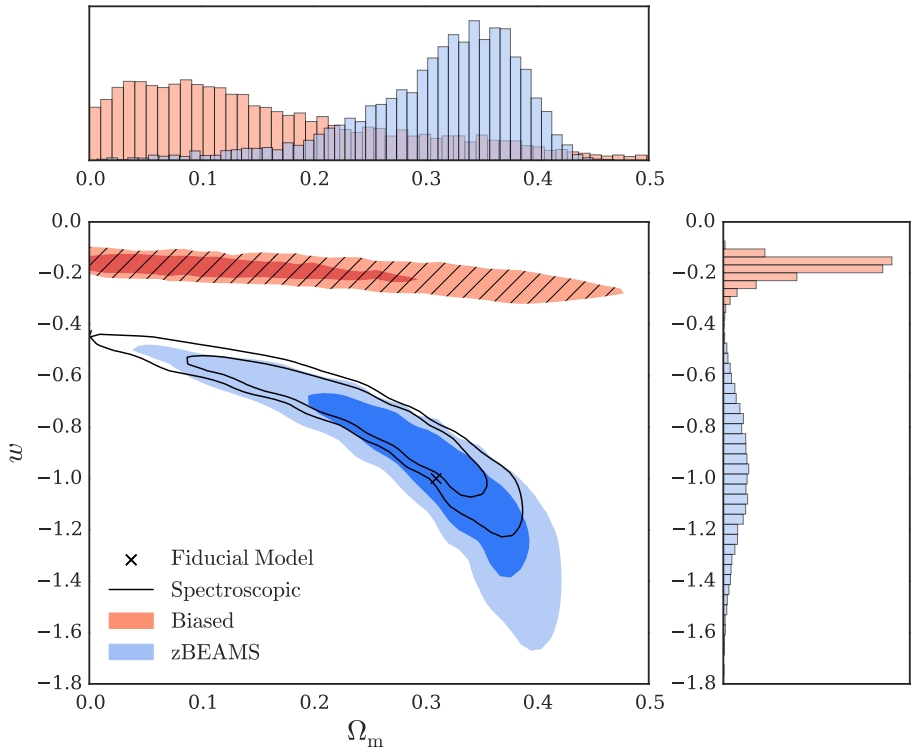

Figure 6. Contour plots for w and Ω_m showing the 68% and 95% credible intervals for the three instances we consider in the photometric case with Gaussian redshift errors $\sigma_z = 0.04(1+z)$. 998 SNIa without type contamination were simulated with redshifts drawn from a prior distribution given by $P(z) \sim ze^{-3z}$. The black cross shows the fiducial model from which the data was generated. The black outlined contours show the ideal "model answer" posterior where we use the standard likelihood on the unbiased/spectroscopic dataset. The red solid contours with dashed lines illustrate the posterior for the standard MCMC on the biased dataset (without taking into account the redshift errors). The blue solid contours show the posterior found using zBEAMS on the same biased dataset. Top and right panels show the 1D marginalised histograms for Ω_m and w respectively for the standard likelihood (red) and the zBEAMS likelihood (blue).

true identity of the galaxy hosting the supernova is unsure, even if the redshifts of potential host galaxies are perfectly known.

In this paper we have shown how to perform unbiased cosmology with such a sample, simultaneously handling both the non-Ia contamination and the problem of imprecise supernova redshifts in a unified framework. Our formalism - zBEAMS - generalises the original BEAMS formalism [25] to handle the redshift uncertainties of the supernovae by employing a hierarchical Bayesian approach. We introduce nuisance parameters for the type and redshift of each supernova in the dataset and then marginalise over these nuisance parameters using numerical sampling. In the special case where the supernova belongs to one of several galaxies each with spectroscopically known redshift, the final posterior is a simple weighted mixture model over the posteriors assuming the supernova is in each of the different potential host galaxies (see eq. (3.13)).

We show in figure 3 that a model with a 9% host misidentification error leads to large bias ($\sim 3\sigma$) using the standard inference approach while zBEAMS removes the bias at essentially no extra computational cost. We also consider the case of photometric uncertainties, where we numerically marginalise over 998 redshift parameters to produce the unbiased contours in figure 6.

There are a number of ways in which this work can be extended in a straight-forward way:

- In this analysis we have assumed that the probability of belonging to a given host galaxy γ, encoded in the terms $P(z|\gamma)$, and the probability of being a given type of supernova, τ, encoded in $P(\tau|z)$, are known *apriori*. It would be interesting to extend our formalism to allow these to be partially known nuisance parameters that are estimated by available data.

- The zBEAMS formalism could be extended to include correlations with host galaxy information, such as host influence on Hubble residuals via stellar mass etc. [31, 32].

- We found that the redshift distribution plays an important role for the photometric redshift case. A much more complex model could be used than the one we assumed which could include systematic effects and allow one to learn something about supernova rates.

- While we have presented the zBEAMS formalism emphasising its generic nature for any data D, in our examples we took D to be the measured distance moduli. It would be useful to develop zBEAMS specifically for the case in which D is the set of lightcurve flux measurements in different bands; i.e. one step further back in the analysis chain.

These extensions are left to future work.

Acknowledgments

The authors thank Alan Heavens, Nosiphiwo Zwane, Laura Richter and other participants of the 2016 Bayes School in Stellenbosch for early contributions and Rahul Biswas, Rick Kessler, Bob Nichol and Hiranya Peiris for comments on the draft. BB thanks Renée Hlozek and Martin Kunz for long-term collaborations related to BEAMS and Roberto Trotta for discussions. ML thanks Brendon Brewer for help with DNest. This paper originated at the 2016 Bayes School and Workshop in November 2016 funded by NITHeP, and grew out of Alan Heaven's lectures at the same School. The authors acknowledge support from the NRF, SKA and AIMS. We acknowledge use of the Hubble ESA Archive for images. This paper and the research it describes were created using scrum methodology.

References

[1] D. H. Weinberg et al., *Observational probes of cosmic acceleration*, *Physics Reports* **530** (2013) 87-255 arXiv:1201.2434

[2] B. A. Bassett and R. Hlozek, *Baryon acoustic oscillations*, *Dark Energy: Observational and Theoretical Approaches* P. Ruiz-Lapuente (2010) 246

[3] LSST Science Collaboration and P. A. Abell, *LSST Science Book, Version 2.0*, *ArXiv e-prints* (2009) arXiv:0912.0201

[4] K. S. Mandel et al., *Type Ia Supernova Light-Curve Inference: Hierarchical Bayesian Analysis in the Near-Infrared*, *Astrophysical Journal* **704** (2009) 629-651 arXiv:0908.0536

[5] D. Rubin et al., *UNITY: Confronting Supernova Cosmology's Statistical and Systematic Uncertainties in a Unified Bayesian Framework*, *Astrophysical Journal* **813** (2015) 137 arXiv:1507.01602

[6] R. Kessler and D. Scolnic, *Correcting Type Ia Supernova Distances for Selection Biases and Contamination in Photometrically Identified Samples*, *Astrophysical Journal* **836** (2017) 56 arXiv:1610.04677

[7] H. Shariff, *BAHAMAS: New Analysis of Type Ia Supernovae Reveals Inconsistencies with Standard Cosmology*, *Astrophysical Journal* **827** (2016) 1 arXiv:1510.05954

[8] C. Ma et al., *Application of Bayesian graphs to SN Ia data analysis and compression*, *MNRAS* **463** (2016) 1651-1665 arXiv:1603.08519

[9] Jennings E. et al., *A new approach for obtaining cosmological constraints from Type Ia Supernovae using Approximate Bayesian Computation*, *ArXiv e-prints* (2016) arXiv:1611.03087

[10] M. Sako et al., *Photometric Type Ia Supernova Candidates from the Three-year SDSS-II SN Survey Data*, *Astrophysical Journal* **738** (2011) 162 arXiv:1107.5106

[11] R. Kessler et al., *Photometric Estimates of Redshifts and Distance Moduli for Type Ia Supernovae*, *Astrophysical Journal* **717** (2010) 40-57 arXiv:1001.0738

[12] Y. Wang et al.,*Analytic photometric redshift estimator for Type Ia supernovae from the Large Synoptic Survey Telescope*, *MNRAS* **451** (2015) 1955-1963 arXiv:1501.06839

[13] A. Möller et al., *Photometric classification of type Ia supernovae in the SuperNova Legacy Survey with supervised learning*, *JCAP* **12** (2016) 008 arXiv:1608.05423

[14] R. Hlozek et al., *Photometric Supernova Cosmology with BEAMS and SDSS-II*, *Astrophysical Journal* **752** (2012) 79 arXiv:1111.5328

[15] H. Campbell et al., *Cosmology with Photometrically Classified Type Ia Supernovae from the SDSS-II Supernova Survey*, *Astrophysical Journal* **763** (2013) 88 arXiv:1211.4480

[16] M. D. Olmstead et al., *Host Galaxy Spectra and Consequences for Supernova Typing from the SDSS SN Survey*, *Astronomical Journal* (2014) 75 arXiv:1308.6818

[17] F. Yuan et al., *OzDES multifibre spectroscopy for the Dark Energy Survey: first-year operation and results*, *MNRAS* **452** (2015) 3047-3063 arXiv:1504.03039

[18] R. R. Gupta et al., *Host Galaxy Identification for Supernova Surveys*, *Astronomical Journal* **152** (2016) 54 arXiv:1604.06138

[19] J. Guy et al., *SALT2: using distant supernovae to improve the use of type Ia supernovae as distance indicators*, *Astronomy and Astrophysics* **466** (2007) 11-21 arXiv:astro-ph/0701828

[20] N. Metropolis et al., *Equation of State Calculations by Fast Computing Machines*, *The Journal of Chemical Physics* **21** (1953) 1087-1092

[21] W. K. Hastings, *Monte Carlo sampling methods using Markov chains and their applications*, *Biometrika* **57** (1970) 97-109

[22] M. C. March et al., *Improved constraints on cosmological parameters from Type Ia supernova data*, *MNRAS* **418** (2011) 2308-2329 arXiv:1102.3237

[23] S. F. Gull, *Bayesian Data Analysis: Straight-line fitting*, *Maximum Entropy and Bayesian Methods* Springer Netherlands (1989) 511-518

[24] A. F. Heavens et al., *Generalized Fisher matrices*, *MNRAS* **445** (2014) 1687-1693 arXiv:1404.2854

[25] M. Kunz et al., *Bayesian estimation applied to multiple species*, *Physical Review D* **75** (2007) 103508 arXiv:astro-ph/0611004

[26] M. Knights et al., *Extending BEAMS to incorporate correlated systematic uncertainties*, *JCAP* **1** (2013) 039 arXiv:1205.3493

[27] R. Kessler et al., *Results from the Supernova Photometric Classification Challenge*, *Publications of the ASP* **122** (2010) 1415 arXiv:1008.1024

[28] Planck collaboration and P. A. R. Ade, *Planck 2015 results. XIII. Cosmological parameters*, *Astronomy and Astrophysics* **594** (2016) A13 arXiv:1502.01589

[29] B. J. Brewer et al., *Diffusive Nested Sampling*, *G. Stat Comput* (2011) 649 arXiv:0912.2380

[30] R. M. Neal, *MCMC using Hamiltonian dynamics*, *ArXiv e-prints* (2012) arXiv:1206.1901

[31] M. Sullivan et al., *The dependence of Type Ia Supernovae luminosities on their host galaxies*, *MNRAS* **406** (2010) 782-802 arXiv:1003.5119

[32] H. Lampeitl et al., *The Effect of Host Galaxies on Type Ia Supernovae in the SDSS-II Supernova Survey*, *Astrophysical Journal* **722** (2010) 566-576 arXiv:1005.4687